

Measuring the non-Gaussian stochastic gravitational-wave background: a method for realistic interferometer data

Eric Thrane^{1, a}

¹*LIGO Laboratory, California Institute of Technology, MS 100-36, Pasadena, CA, 91125, USA*

A stochastic gravitational-wave background (SGWB) can arise from the superposition of many independent events. If the rate of events per unit time is sufficiently high, the resulting background is Gaussian, which is to say that it is characterized only by a gravitational-wave strain power spectrum. Alternatively, if the event rate is low, we expect a non-Gaussian background, characterized by intermittent sub-threshold bursts. Many experimentally accessible models of the SGWB, such as the SGWB arising from compact binary coalescences, are expected to be of this non-Gaussian variety. Primordial backgrounds from the early universe, on the other hand, are more likely to be Gaussian. Measuring the Gaussianity of the SGWB can therefore provide additional information about its origin. In this paper we introduce a novel maximum likelihood estimator that can be used to estimate the non-Gaussian component of an SGWB signature measured in a network of interferometers. This method can be robustly applied to spatially separated interferometers with colored, non-Gaussian noise. Furthermore, it can be cast as a generalization of the widely used stochastic radiometer algorithm.

PACS numbers: 95.55.Ym

I. INTRODUCTION

A stochastic gravitational-wave background (SGWB) is expected to arise from the superposition of many systems, which are individually too weak to detect, but which combine to produce a GW signature characterized by its ensemble statistical properties. In astrophysical models [1], an SGWB can arise from objects such as compact binary coalescences [2–6], neutron stars (including highly magnetized neutron stars) [7–12], young or spun-up neutron stars [9, 13–16], core collapse supernovae [17–20], and white dwarf binaries [21]. Cosmological/primordial sources, meanwhile, can arise from inflationary physics [22–26], cosmic strings [27–32], and pre-Big-Bang models [33, 34].

The initial LIGO and Virgo experiments have yielded a number of constraints on the SGWB [35–37] including limits on the energy density of GWs, which surpass indirect bounds from Big Bang nucleosynthesis and measurements of the cosmic microwave background. A worldwide network of second-generation GW interferometers are expected to begin taking data in 2015 [38–42], and recent work [5] indicates that realistic astrophysical models can be probed with second-generation advanced interferometers—most notably, the SGWB arising from binary neutron star and binary black hole coalescences.

In the event of a detection, it may not be immediately clear *which* systems give rise to the observed SGWB. The strain power spectrum provides one tool for disentangling different possible sources [43, 44]. Also, measurements of the SGWB can be compared with measurements of GW transients in order to indirectly infer information about the SGWB [43]. Finally, sky maps of GW power and tests

of isotropy provide yet another means of characterizing different models [32, 37, 45].

In this paper we explore the *Gaussianity* of the SGWB. A Gaussian SGWB is described only by its strain power spectrum, while a non-Gaussian SGWB (sometimes referred to as an SGWB in the “popcorn” or “shot noise” regime) consists of a series of discrete sub-threshold bursts (see Figure 1). Non-Gaussian signals can be described with a probability distribution of burst waveforms and a duty cycle ξ , which we define as the fraction of data segments during which a GW source somewhere in the universe emits GWs in some analysis band. In the analysis that follows, it will be useful to divide the full GW observing band into smaller analysis bands.

In this work we assume the duty cycle is less than unity, i.e., it is rare for two or more events to simultaneously emit in the same band. For many astrophysical models, this is a good approximation. In the case of the SGWB from binary neutron star coalescence, for example, we expect that $\xi \approx 0.5\%$ for a 4 Hz-wide bin centered at 100 Hz. We refer to both Gaussian and non-Gaussian signals as “stochastic” since both can be described in terms of the ensemble behavior of many individually undetectable bursts. For a more nuanced discussion of relevant terminology, see Section A in the appendix.

Different models of the SGWB predict different levels of Gaussianity (see, e.g., [46]). Early-universe scenarios, as a rule of thumb, produce nearly Gaussian signatures whereas astrophysical scenarios tend to be more non-Gaussian [47]. Moreover, a single model can produce a range of different duty cycles and burst amplitudes depending on its parameters. Measurements of SGWB Gaussianity can provide an important probe to distinguish between models and also to estimate cosmological parameters such as the GW burst rate.

We introduce a maximum likelihood statistic to estimate the duty cycle and other parameters associated

^aElectronic address: ethrane@ligo.caltech.edu

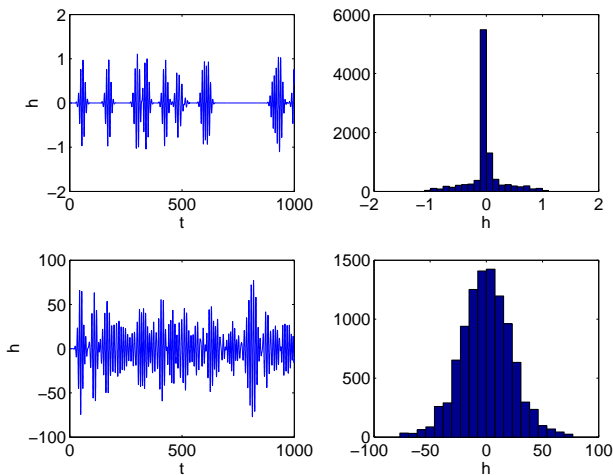


FIG. 1: An illustration of Gaussian and non-Gaussian signals. Top-left: a time series of sine-Gaussian bursts with a low duty cycle produces a non-Gaussian signal. Top-right: histogram of this non-Gaussian signal. Bottom-left: a time series of sine-Gaussian bursts with a high duty cycle produces an approximately Gaussian signal. Bottom-right: histogram of this Gaussian signal.

with non-Gaussian SGWB signatures in GW interferometers. In II C, we show that this statistic can be cast as a generalization of the stochastic radiometer [48, 49]. Previous work has yielded a nearly optimal technique (in the statistical sense) for the special case of colocated, co-aligned interferometers characterized by stationary, Gaussian white noise [50]. More recently, Seto has proposed the use of higher order moments for measuring the Gaussianity of the SGWB [51, 52].

Real GW interferometer data, however, is far from idealized noise. It is colored, non-Gaussian, non-stationary, and colocated interferometers suffer from correlated noise. These factors have made it challenging to implement the method from [50] in practice. The maximum likelihood method proposed here is applicable to realistic GW interferometer noise.

In Section II, we review the canonical framework for analyzing the Gaussian SGWB and introduce a new formalism for handling non-Gaussian signals; in Section III, we demonstrate the non-Gaussian formalism with a Monte Carlo simulation; in Section IV we discuss the sensitivity of SGWB searches using the non-Gaussian formalism. Finally, in Section V, we summarize our findings and describe the next steps necessary to implement the non-Gaussian formalism with actual data and for specific models.

II. FORMALISM

A. Gaussian searches

In order to motivate our non-Gaussian statistic, we begin by describing how traditional (Gaussian) SGWB analyses are carried out. The basic idea behind a Gaussian SGWB measurement is to cross-correlate strain time series from a pair of detectors I and J to create a cross correlation statistic, which is an estimator for the energy density spectrum of GWs (see, e.g., [35, 53]):

$$\Omega_{\text{GW}}(f) = \frac{f}{\rho_c} \frac{d\rho_{\text{GW}}}{df}. \quad (2.1)$$

Here f is the GW frequency, $d\rho_{\text{GW}}$ is the energy density of GWs in a frequency band $(f, f + df)$, and ρ_c is the critical energy density of the universe. By summing data from many time segments and frequency bins, it is possible to observe signals orders of magnitude smaller than the instantaneous noise curve. For additional details, the interested reader is referred to [53].

The SGWB is assumed to be isotropic, stationary, and Gaussian [53, 54]. The measured strain time series in detector I can be written as

$$s_I(t) = h_I(t) + n_I(t). \quad (2.2)$$

Here $h_I(t)$ and $n_I(t)$ are respectively the astrophysical strain signal and the strain noise. For spatially separated interferometers, $n_I(t)$ and $n_J(t)$ are expected to be uncorrelated. Signal and noise are, of course, expected to be uncorrelated. The signals in I and J , however, are expected to be highly correlated:

$$\langle \tilde{h}_I^*(t; f) \tilde{h}_J(t; f) \rangle \neq 0. \quad (2.3)$$

Here, tildes denote discrete Fourier transforms, and $(t; f)$ are spectrogram indices for time and frequency bins respectively.

From Eq. 2.2-2.3, we can construct an estimator from the cross-power spectrum of the I and J strain channels [53]:

$$\hat{Y}(t; f) = \frac{2}{T} Q(f) \text{Re}[\tilde{s}_I^*(t; f) \tilde{s}_J(t; f)]. \quad (2.4)$$

T is the segment duration and $Q(f)$ is a filter function,

$$Q(f) \equiv \frac{f^3}{\gamma(f) \Omega_M(f)} \quad (2.5)$$

which emphasizes certain frequency bins depending on the spectral shape $\Omega_M(f)$ of a particular model, and a geometric factor called the overlap reduction function $\gamma(f)$ (see Figure 2), which takes into account the fact that uncorrelated GW signals from different parts of the sky interfere to reduce the observed correlation between spatially separated interferometers.

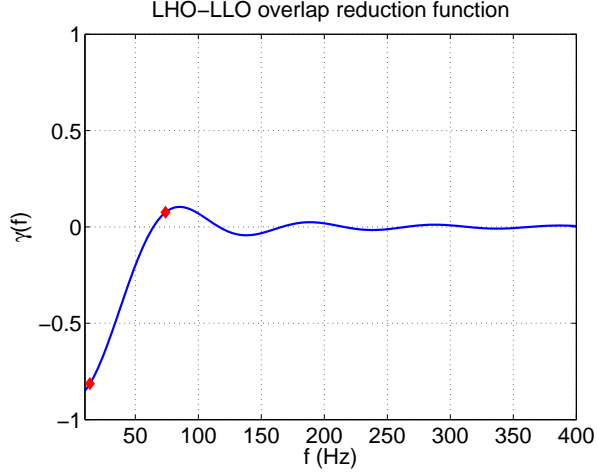


FIG. 2: The overlap reduction function $\gamma(f)$ for the LIGO Hanford - LIGO Livingston detector pair. The red diamonds mark frequency bins of $f = 14$ Hz and 74 Hz, which are singled out below to illustrate how the distribution of non-Gaussian signal varies with frequency.

An estimator for the variance of $\hat{Y}(t; f)$ is given by

$$\hat{\sigma}^2(t; f) = \frac{|Q(f)|^2}{2T\delta f} P_I(t; f) P_J(t; f), \quad (2.6)$$

where δf is the frequency resolution and $P_I(t; f)$ is the auto-power in detector I in time-frequency bins $(t; f)$ [65]. Armed with $\hat{Y}(t; f)$ and $\hat{\sigma}(t; f)$, the optimal estimator for $\int df \Omega_{\text{GW}}(f)$ at time t is given by a weighted average:

$$\begin{aligned} \hat{Y}(t) &= \sum_f \hat{Y}(t; f) \hat{\sigma}^{-2}(t; f) / \sum_f \sigma^{-2}(t; f) \\ \hat{\sigma}^{-2}(t) &= \sum_f \hat{\sigma}^{-2}(t; f). \end{aligned} \quad (2.7)$$

The optimal estimator for the entire data-taking period is a weighted average over time:

$$\begin{aligned} \hat{Y}_{\text{tot}} &= \frac{\sum_t \hat{Y}(t) \hat{\sigma}^{-2}(t)}{\sum_t \hat{\sigma}^{-2}(t)} \\ \hat{\sigma}_{\text{tot}}^{-2} &= \sum_t \hat{\sigma}^{-2}(t) \\ \text{SNR}_{\text{tot}} &= \hat{Y}_{\text{tot}} / \hat{\sigma}_{\text{tot}}. \end{aligned} \quad (2.8)$$

SNR_{tot} is expected to be well-approximated by a normal distribution by the central limit theorem, and indeed, this is born out empirically [35, 37].

B. A non-Gaussian statistic

A number of applications have emerged that make use of $\hat{Y}(t; f)$ and $\hat{\sigma}(t; f)$ —the intermediate data products

that go into the calculation of \hat{Y}_{tot} and $\hat{\sigma}_{\text{tot}}$ [55–57]. Recent work utilizes spectrograms of cross- and auto-power to search for long-lived GW transients and to identify sources of environmental noise contaminating strain data channels. Building on this work, we cast our non-Gaussian statistic in terms of these intermediate data.

To begin, we define a complex estimator

$$\hat{Y}'(t; f) = Q(f) \hat{s}_I^*(t; f) \hat{s}_J(t; f), \quad (2.9)$$

which is a simple generalization of Eq. 2.4. Unlike a Gaussian background, which is isotropic at every instant in time, non-Gaussian bursts are associated with individual sky locations, (even if, on average, they are drawn from an isotropic distribution). This information is encoded in the phase of $\hat{Y}'(t; f)$, and so it is necessary to work with both real and imaginary components.

We take as our starting point $\hat{Y}'(t; f)$ and $\hat{\sigma}(t; f)$ and their ratio

$$\rho(t; f) = \hat{Y}'(t; f) / \hat{\sigma}(t; f). \quad (2.10)$$

The quantity $\rho(t; f)$ has useful properties for our purposes. First, it is well-studied and already in use in stochastic analyses. Second, its statistical behavior can be probed robustly through time slides in which one strain time series is offset by an amount greater than the GW travel time between detectors in order to obtain many independent realizations of noise.

In order to derive our non-Gaussianity statistic, we endeavor to answer a simple question: how does the signal distribution of $\rho(t; f)$, denoted S , differ from the background distribution of $\rho(t; f)$, denoted B . We take each value of $(t; f)$ —pixels in spectrograms of $\rho(t; f)$ —to be a separate measurement.

For our purposes, *any measurement in which there is a non-Gaussian burst signal in $(t; f)$ is drawn from S* . All other measurements are considered to be background. This definition of signal and background is useful, but it can be counterintuitive. The signal distribution, as we have defined it, is determined not only by properties of the non-Gaussian SGWB; it is also determined in part by properties of the detector noise. Defining signal and background like this will be useful to derive an estimator for duty cycle.

Having pointed out these subtleties, we can now write the signal distribution as

$$S(\rho(t; f) | \vec{\theta}_{\text{signal}}, \vec{\theta}_{\text{noise}}), \quad (2.11)$$

where $\vec{\theta}_{\text{signal}}$ is a vector of parameters describing the SGWB and $\vec{\theta}_{\text{noise}}$ describes the detector noise. In the same vein, we can write the background distribution as

$$B(\rho(t; f) | \vec{\theta}_{\text{noise}}). \quad (2.12)$$

As we proceed we shall refer to simply $S(\vec{\Theta})$ and $B(\vec{\Theta})$, using capital $\vec{\Theta}$ as shorthand for the appropriate vector of parameters. Further, we drop $(t; f)$ arguments in favor

of i , which denotes separate measurements. We return later to asymmetries in the t and f variables.

Examples of $S(\rho_i|\vec{\Theta})$ and $B(\rho_i|\vec{\Theta})$ obtained using Monte Carlo are shown in Figure 3. A detailed description of the Monte Carlo simulation is provided in the subsequent section. For now, we simply describe the panels of Figure 3 in general terms and point out some of the interesting features.

The top-left panel shows a scatter plot of ρ for Gaussian noise. (In this case, we consider a frequency bin centered on $f = 74$ Hz, but the distribution looks the same at all frequencies.) The distribution has a mean of zero and is narrow compared to both non-Gaussian signal and non-Gaussian noise. The top-right panel, meanwhile, shows the case of a non-Gaussian signal in the presence of Gaussian noise in a frequency bin centered on $f = 14$ Hz. At 14 Hz, the GW wavelength $\lambda \approx 2 \times 10^4$ km is large compared to the separation of the interferometers, ≈ 3000 km for the case of the LIGO Hanford Observatory (LHO) and LIGO Livingston Observatory (LLO) as used in this simulation. Thus, ρ is distributed approximately as it would be for a colocated pair of interferometers. The presence of a signal is evidenced by the shift of the distribution away from zero. The fact that the shift is negative is due to relative orientations of the LHO and LLO detectors, which is encoded in the sign of the overlap reduction function (see Figure 2).

The lower-left panel shows the distribution of ρ for non-Gaussian signal in the presence of Gaussian noise in a frequency bin centered on $f = 74$ Hz. At 74 Hz, the interferometers no longer behave as though they are colocated. The presence of a signal shifts the mean away from zero, but it also changes the width and shape of the distribution. Since 74 Hz occurs in between the first and second zeros of the overlap reduction function, the mean is positive (see Figure 2). For illustrative purposes, the non-Gaussian bursts used to make this plot are made loud compared to the detector noise.

The lower-right panel shows non-Gaussian detector noise (also called “glitchy” noise) simulated by taking the same bursts used to simulate a non-Gaussian signal (lower-left), but generating a signal in only one detector at a time; i.e., we assume the noise at each site is uncorrelated. The glitchiness of the signal widens the distribution compared to Gaussian noise, but not as much as coincident non-Gaussian signals, which suggests that we can differentiate population of glitches from non-Gaussian signals. Additionally, the mean of the non-Gaussian noise distribution is still zero, like Gaussian noise. By comparing these distributions of different signal and noise distributions, we can build a statistical framework to differentiate them. This is the goal of the remainder of this section.

In practice, $S(\rho_i|\vec{\Theta})$ and $B(\rho_i|\vec{\Theta})$ can be calculated from Monte Carlo or from pseudo experiments performed with time-shifted data. The former method is computationally cheaper, but yields a less accurate description of the noise. Even an approximate Monte Carlo descrip-

tion, however, can be a useful starting point. $S(\rho_i|\vec{\Theta})$ and $B(\rho_i|\vec{\Theta})$ are used to weight data as more or less like signal. To the extent that they differ from the true distributions, the likelihood statistic will be less effective distinguishing between signal and background. However, if the likelihood statistic is ultimately tested empirically on time-shifted data, then we can avoid bias in detection or parameter estimation—even if we construct the estimator with an only approximate Monte Carlo model.

Armed with distributions of signal and background, the likelihood of observing ρ_i can be written as

$$P_i(\xi|\rho_i, \vec{\Theta}) = \xi S(\rho_i|\vec{\Theta}) + (1 - \xi)B(\rho_i|\vec{\Theta}). \quad (2.13)$$

Here ξ is the probability that the measurement is drawn from the signal distribution. We can interpret ξ as the duty cycle for our non-Gaussian signal model. This formulation is similar to a maximum likelihood approach used in neutrino and gamma-ray astronomy, see, e.g., [58–60].

Given N measurements, we can construct a likelihood function

$$\mathcal{L}(\xi|\{\rho_i\}, \vec{\Theta}) = \prod_i^N P_i(\rho_i|\rho_i, \vec{\Theta}). \quad (2.14)$$

Here, we have implicitly assumed that each measurement is statistically independent. This assumption is expected to be approximately valid for reasonably stationary noise and for signals that do not repeatedly wander in and out of the same frequency bin. Subtle effects from overlapping data segments may require a more careful treatment. An estimator for ξ (denoted $\hat{\xi}$) is given simply by maximizing $\mathcal{L}(\xi)$. Confidence intervals for ξ are calculated straightforwardly from $\mathcal{L}(\xi|\{\rho_i\}, \vec{\Theta})$. To illustrate, we perform a simple Monte Carlo calculation, described in detail in Section III and summarized in Figure 4.

Finally, we consider the question of how best to detect non-Gaussianity in an SGWB signal. We can frame this question more precisely as: what is the appropriate metric for determining whether ξ is non-zero. A convenient measure of non-Gaussian signal strength is the ratio of the likelihood evaluated at its maximum to the likelihood evaluated at $\xi = 0$ (see [58–60]):

$$\Lambda = 2 \log \left[\mathcal{L}(\hat{\xi}) / \mathcal{L}(0) \right]. \quad (2.15)$$

If $S(\rho_i|\vec{\Theta})$ and $B(\rho_i|\vec{\Theta})$ are accurate descriptions of the signal and background, then the probabilistic interpretation of Λ as (twice the log of) a likelihood ratio is straightforward. However, even if $S(\rho_i|\vec{\Theta})$ and $B(\rho_i|\vec{\Theta})$ are only approximately known, we can calculate Λ for many realizations of time-shifted data. In this way, we can obtain a robust and empirical means of converting Λ to a false alarm probability—even if our noise and signal models are only roughly approximate.

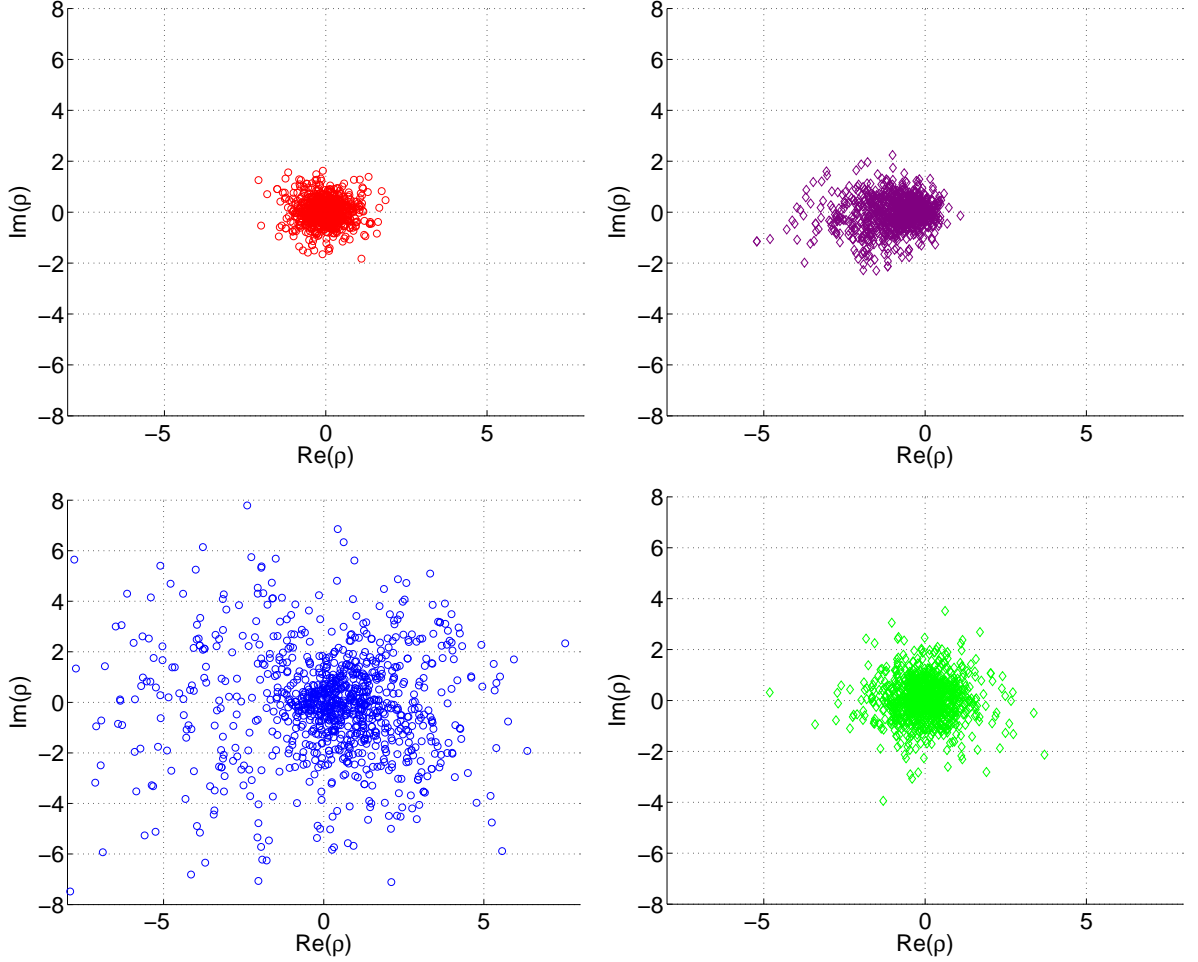


FIG. 3: Scatter plots of ρ for different signal and background models. Top-left: Gaussian noise with no signal (arbitrary frequency band). Top-right: Gaussian noise in the presence of a non-Gaussian signal; ($f = 14$ Hz). Bottom-left: Gaussian noise in the presence of a non-Gaussian signal ($f = 74$ Hz). Bottom-right: Gaussian noise with non-Gaussian glitches identical to the non-Gaussian signals, but only in one detector during each measurement ($f = 74$ Hz). In each plot, the frequency bin width is 4 Hz.

C. Relationship to radiometer statistic

The stochastic radiometer statistic [48, 49] applies the Gaussian isotropic formalism of [53] to the case of a Gaussian point source. In this subsection we show how the radiometer statistic can be cast as a special case of our non-Gaussian statistic. We begin by defining the signal model. We consider a point source associated with a fixed sky location \hat{n} . We assume that the source is characterized by a stationary GW energy density spectrum $\Omega(f) = \bar{\Omega}$. (For the sake of simplicity, we assume that it is constant with respect to f .)

For a single point source, there is a known phase relationship between I and J and so Eq. 2.4 becomes [48]

$$\hat{Y}(t; f) = Q(f) \text{Re} \left[e^{-2\pi i f \tau(\hat{n}, t)} \tilde{s}_I^*(t; f) \tilde{s}_J(t; f) \right]. \quad (2.16)$$

Here $\tau(\hat{n}, t)$ is the delay time between the interferometers.

The signal distribution is given approximately by [55]

$$S(\rho_i | \bar{\Omega}) \propto \exp \left(-\frac{1}{2} (\rho_i - \bar{\Omega} / \sigma_i)^2 \right). \quad (2.17)$$

Since the source is, by assumption, always emitting GWs, we can set $\xi = 1$, which implies that there is GW signal present in every data segment and that we can ignore $B(\rho_i | \bar{\Omega})$ entirely. It follows from Eqs. 2.13-2.14 that the likelihood can be written as

$$\begin{aligned} \mathcal{L}(\xi = 1 | \{\rho_i\}, \bar{\Omega}) &\propto \exp \left(-\frac{1}{2} \sum_i^N (\rho_i - \bar{\Omega} / \sigma_i)^2 \right) \\ &\propto \exp \left(-(\bar{\Omega} - \hat{Y}_{\text{tot}})^2 / 2\sigma_{\text{tot}}^2 \right), \end{aligned} \quad (2.18)$$

where \hat{Y}_{tot} is simply the optimal estimator from Eq. 2.8.

Thus, the radiometer statistic is a special case of the non-Gaussian statistic in which the signal model fixes the

sky location of the source and the duty cycle is set to $\xi = 1$. The likelihood function can then be used to estimate the energy density spectrum $\bar{\Omega}$. In principle, a similar analogy is possible between the isotropic statistic [53] and the non-Gaussian statistic. However, the fact that both the radiometer and the non-Gaussian statistic assume GW point sources makes the analogy shown here more straightforward.

D. Comparison to other methods

One of the first papers to address the topic of detecting a non-Gaussian SGWB is [50]. Our method differs in several important ways. In this work we relax the assumptions from [50] that the noise is Gaussian and white. Instead of relying on two colocated detectors as in [50], we assume two spatially separate detectors. Unlike [50], our likelihood statistic is framed in terms of cross-power, with auto-power terms used only for normalization.

Since the statistic in [50] is nearly optimal, it is very likely to provide a more sensitive measurement within its domain of utility, compared to the method described here. However, our statistic (built from cross-power) can be extended straightforwardly to interferometers with colored, non-Gaussian noise, and they need not be colocated. Spatially separated interferometers, in turn, allow for the use of robust background estimation through time slides.

Statistics utilizing forth-order (kurtosis) strain moments have been proposed as probes for non-Gaussianity in the SGWB [51, 52]. It is presently difficult to evaluate the relative merits of different non-Gaussianity techniques, since none has been utilized with real interferometer data. Clearly, the next step is to carry out analyses with real data.

III. SIMULATION AND RESULTS

In order to demonstrate our likelihood formalism, we perform a Monte Carlo simulation. We generate three types of data: Gaussian noise, Gaussian noise + non-Gaussian GW bursts (signal), and non-Gaussian noise. The Gaussian noise is colored according to the design sensitivity of Advanced LIGO [38] and we assume a network consisting of 4 km detectors at LHO and LLO. We employ a toy signal model consisting of randomly arriving 200 ms white-noise bursts with a strain amplitude density of $\approx 3 \times 10^{-24} \text{ Hz}^{-1/2}$. These bursts are marginal compared to the noise—the average $|\hat{\rho}|$ is 0.43 for bursts plus noise and only slightly less (0.41) for noise alone—and can therefore be characterized as sub-threshold.

We calculate spectrograms of $\rho(t; f)$ (Eq. 2.10) using $4 \text{ Hz} \times 1 \text{ s}$ pixels. Since we are presently concerned only with demonstration, this choice of pixel size is arbitrary. The issue of pixel size is revisited in the appendix.

In order to construct the likelihood function used in Figure 4, we construct distributions of $S(\rho_i|\vec{\Theta})$ and $B(\rho_i|\vec{\Theta})$ using 10^7 trials of Monte Carlo data. We then use an independent dataset consisting of 500 Monte Carlo measurements, which—following Eq. 2.14—we compare to $S(\rho_i|\vec{\Theta})$ and $B(\rho_i|\vec{\Theta})$ in order to measure the duty cycle ξ [66].

The results of our simulation are illustrated in Figure 4. We test three datasets: one consisting of pure background ($\xi = 0$, dash-dot blue), one consisting of pure signal ($\xi = 1$, solid red), and a third consisting of a 50% mixture of each ($\xi = 0.5$, dashed green). In all three cases we find that the observed posterior distributions are consistent with the true value of ξ . This demonstrates, with a very simple toy model, how our formalism can be used to measure ξ in the presence of colored noise in spatially separated interferometers.

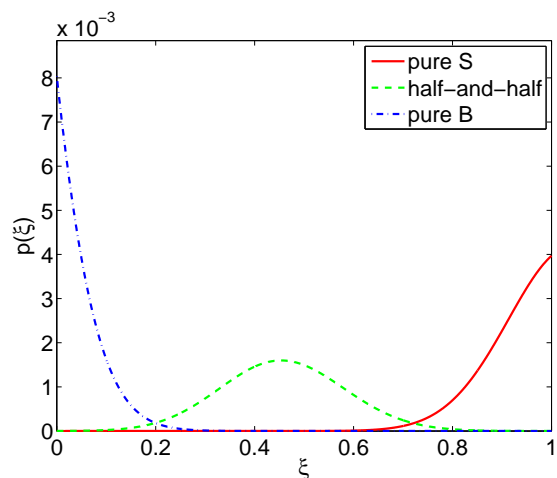


FIG. 4: Example posteriors for duty cycle ξ using Monte Carlo data for pure background $\xi = 0$ (dash-dot blue), pure signal $\xi = 1$ (solid green), and an even mixture of the two $\xi = 0.5$ (dashed green).

IV. SENSITIVITY

A natural question is: if the non-Gaussian search presented here incorporates information about the non-Gaussian character of the popcorn signal we seek to measure, can it, in some cases, provide a more sensitive detection statistic than the Gaussian statistic used in previous stochastic searches [35–37]? A detailed analysis, beyond our present scope, is required to answer this question thoroughly. However, there are several points worth noting.

First, in the limit of (highly idealized) stationary Gaussian noise, we expect the non-Gaussian statistic will outperform the Gaussian statistic. To illustrate, we note that the green data in Figure 4 have a Gaussian statistic signature of $\text{SNR}_{\text{tot}} < 1$ (typical of pure noise) whereas

the non-Gaussian statistic $\Lambda = 15$ represents a strong detection. We also expect, however, that as the data becomes glitchier, the advantage of the non-Gaussian approach will diminish, since both glitches and non-Gaussian bursts will have a tendency to perturb higher-order moments of the distribution of $\rho(t; f)$ (albeit in different ways).

Second, the Gaussian statistic is almost completely insensitive to stochastic signals in frequency bins corresponding to the zeros of the overlap reduction function $\gamma(f)$. These zeros represent frequencies at which the detector pair are as likely to be out of phase as in phase, and so the integrated signal is zero. Since the non-Gaussian technique presented here incorporates higher-order moments in the distribution of $\rho(t; f)$ (beyond the mean), it will have at least some sensitivity to the SGWB even when $\gamma(f) \approx 0$.

Third, while there are potential advantages associated with the non-Gaussian statistic, it is worthwhile to mention several advantages possessed by the Gaussian statistic. It is very well studied and has been shown to yield reliable results [35–37, 53, 61–63], it is simple to understand and implement, and since it utilizes the sum of a great many numbers, it is very robust to non-stationary noise artifacts. Thus, the Gaussian statistic is likely to provide an important benchmark and cross-check to results obtained with the non-Gaussian statistic.

V. CONCLUSIONS

Some of the most promising sources of the stochastic gravitational-wave background (such as compact binary coalescences) are likely to be non-Gaussian. By measuring the non-Gaussianity of the stochastic background, we can learn more about its origin. To this end, we have presented a maximum likelihood estimator that can be used to measure the non-Gaussianity of the stochastic gravitational-wave background utilizing realistic interferometer data. Using Monte Carlo data, we illustrated how the calculation can be carried out, and demonstrated that we can estimate the duty cycle of the bursts that characterize a non-Gaussian signal. We outlined the next steps, which must be undertaken in order to tune the analysis for specific astrophysical models such as the stochastic background arising from compact binary coalescences. Future work will focus on carrying out this optimization.

Acknowledgments

We thank Joseph Romano and Tania Regimbau for thoughtful comments on a draft of this paper. This work is by a member of the LIGO Laboratory, supported by funding from United States National Science Foundation. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Founda-

tion and operates under cooperative agreement PHY-0757058. This paper has been assigned LIGO document number ligo-p1200141.

Appendix A: Terminology

Rosado [47] has reviewed the SGWB literature and attempted to provide a comprehensive and standardized glossary of terminology. Where possible, we try to follow the terminology of [47], referring, for example, to the objects, which combine to create a SGWB as “systems.” However, while we are loath to add to the SGWB lexicon, some definitions and distinctions are necessary for our present purpose.

Rosado makes a distinction between (in-principle) *resolvable* and *unresolvable* SGWBs. Unresolvable SGWB signals, according to [47], are present when, on average, two or more systems simultaneously create strain signals in the same frequency bin. This distinction is most useful in the context of a far-future detector with sufficient sensitivity to subtract out resolvable signals in order to measure an underlying primordial SGWB [64]. Near-future detectors will lack the sensitivity to separately measure the systems contributing to the “resolvable” SGWB. It is therefore useful to define *sub-threshold bursts* as the components of an in-principle resolvable SGWB, which cannot be resolved in practice. An SGWB consisting of sub-threshold bursts is always resolvable according to the definition in [47]. Whether or not a resolvable SGWB consists of sub-threshold bursts will depend on the detector used to measure it. The non-Gaussian SGWB considered in this paper consist of sub-threshold bursts.

Appendix B: Details

Here we point out details that will require more careful attention in order to implement this method for a specific SGWB model. Our aim is not to provide a systematic treatment, but rather to highlight some of the finer points worthy of attention.

Probability density functions.—The distributions for $S(\rho_i|\vec{\Theta})$ and $B(\rho_i|\vec{\Theta})$ must be sampled with sufficient resolution to distinguish between signal and background. Models with very low-level bursts may require very high resolutions, and so significant computational resources may be necessary to compute $S(\rho_i|\vec{\Theta})$ and $B(\rho_i|\vec{\Theta})$.

Pixel size.—Pixel size can be chosen to optimize the sensitivity of a search. The pixel dimensions should be chosen so as to be comparable to the time and bandwidth of the non-Gaussian burst that is the target of the search. Pixels that are very long/short in time are undesirable because, in the first case, the signal will be diluted with more noise than necessary, and in the second case, the signal will be spread thinly over many pixels. An analogous argument can be made for the frequency bin

width. Numerical studies determine a suitable pixel size appropriate for a given model.

Broadband analysis.—The behavior of $S(\rho_i|\vec{\Theta})$ and $B(\rho_i|\vec{\Theta})$ varies significantly depending on the frequency band of interest. For example, the mean of $S(\rho_i|\vec{\Theta})$ can be positive, negative, or zero depending on the value of the overlap reduction function at the frequency in question (see Figure 2). (For bursts drawn from an isotropic distribution, $\langle\rho_i\rangle$ is always real since ρ 's drawn from some direction \hat{n} have, on average, the opposite imaginary component of ρ 's drawn from the antipodal direction $-\hat{n}$.) Therefore, it may be desirable to calculate $S(\rho_i|\vec{\Theta})$ and $B(\rho_i|\vec{\Theta})$ for many different bands.

Additional complications may arise from the fact that the signal may not spend the same duration emitting in every band. Compact binary coalescences, for example, emit at a frequency that accelerates as a function of time. In order to combine posterior distributions of

ξ from different bands, it may therefore be necessary to apply normalization factors to take into account the expected frequency evolution of the signal.

Time-varying detector performance.—For a variety of reasons, real GW detectors vary in performance on timescales ranging from minutes to months. As examples, anthropogenic noise can cause elevated noise levels during the local rush hour, and noise performance can improve month to month following commissioning breaks. In the example plots showed in Figure 3, we assume that the detector noise is stationary.

In order take into account the variability of the noise as a function of time, it may prove useful to add another variable to $\vec{\Theta}$ describing the variability of the noise. Another option could be to simply use a subset of the highest quality data in which the strain sensitivity and glitchiness are relatively uniform.

-
- [1] T. Regimbau, Res. Astr. Astrop. **11**, 369 (2011).
 - [2] E. S. Phinney, Astrophys. J. Lett. **380**, L17 (1991).
 - [3] D. I. Kosenko and K. A. Postnov, Astron. & Astrop. **336**, 786 (1998).
 - [4] X.-J. Zhu, E. Howell, T. Regimbau, D. Blair, and Z.-H. Zhu, Astrophys. J. **739**, 86 (2011).
 - [5] C. Wu, V. Mandic, and T. Regimbau, Phys. Rev. D **85**, 104024 (2012).
 - [6] T. Regimbau and J. A. de Freitas Pacheco, Astrophys. J. **642**, 455 (2006).
 - [7] C. Cutler, Phys. Rev. D **66**, 084025 (2002).
 - [8] T. Regimbau and V. Mandic, Class. Quantum Grav. **25**, 184018 (2008).
 - [9] T. Regimbau and J. A. de Freitas Pacheco, Astron. and Astrophys. **376**, 381 (2001).
 - [10] P. A. Rosado, Phys. Rev. D **86**, 104007 (2012).
 - [11] E. Howell, T. Regimbau, A. Corsi, D. Coward, and R. Burman, Mon. Not. R. Ast. Soc. **410**, 2213 (2010).
 - [12] S. Marassi, R. Ciolfi, R. Schneider, L. Stella, and V. Ferrari, **411**, 2549 (2011).
 - [13] B. J. Owen, L. Lindblom, C. Cutler, B. F. Schutz, A. Vecchio, and N. Andersson, Phys. Rev. D **58**, 084020 (1998).
 - [14] S. Chandrasekhar, *Ellipsoidal Figures of Equilibrium* (New Haven, Yale Univ. Press, 1969).
 - [15] J. L. Houser, J. M. Centrella, and S. C. Smith, Phys. Rev. Lett. **72**, 1314 (1994).
 - [16] D. Lai and S. L. Shapiro, Astrophys. J. **442**, 259 (1995).
 - [17] P. Sandick, K. A. Olive, F. Daigne, and E. Vangioni, Phys. Rev. D **73**, 104024 (2006).
 - [18] E. Howell, D. Coward, R. Burman, D. Blair, and J. Gilmore, Mon. Not. R. Ast. Soc. **351**, 1237 (2004).
 - [19] A. Buonanno and G. Sigl and G. G. Raffelt and H.-T. Janka and E. Müller, Phys. Rev. D **72**, 084001 (2005).
 - [20] S. Marassi, R. Schneider, and V. Ferrari, Mon. Not. R. Ast. Soc. **398**, 293 (2009).
 - [21] A. J. Farmer and E. Phinney, Mon. Not. R. Ast. Soc. **346**, 1197 (2003).
 - [22] L. P. Grishchuk, Sov. Phys. JETP **40**, 409 (1975).
 - [23] A. A. Starobinskii, JETP Lett. **30**, 682 (1979).
 - [24] R. Easther and E. A. Lim, JCAP **0604**, 010 (2006).
 - [25] N. Barnaby, E. Pajer, and M. Peloso, Phys. Rev. D **85**, 023535 (2012).
 - [26] J. Cook and L. Sorbo, Phys. Rev. D **85**, 023534 (2012).
 - [27] R. R. Caldwell and B. Allen, Phys. Rev. D **45**, 3447 (1992).
 - [28] T. Damour and A. Vilenkin, Phys. Rev. Lett. **85**, 3761 (2000).
 - [29] T. Damour and A. Vilenkin, Phys. Rev. D **71**, 063510 (2005).
 - [30] X. Siemens, V. Mandic, and J. Creighton, Phys. Rev. Lett. **98**, 111101 (2007).
 - [31] S. Olmez, V. Mandic, and X. Siemens, Phys. Rev. D **81**, 104028 (2010).
 - [32] S. Olmez, V. Mandic, and X. Siemens, arXiv:1106.5555 (2011).
 - [33] M. Gasperini and G. Veneziano, Astropart. Phys. **1**, 317 (1993).
 - [34] V. Mandic and A. Buonanno, Phys. Rev. D **73**, 063008 (2006).
 - [35] B. Abbott et al., Nature **460**, 990 (2009).
 - [36] J. Abadie et al., Phys. Rev. D **85**, 122001 (2012).
 - [37] J. Abadie et al. (LIGO Scientific Collaboration, Virgo Collaboration), Phys. Rev. Lett. **107**, 271102 (2011).
 - [38] G. M. Harry (for the LIGO Scientific Collaboration), Class. Quantum Grav. **27**, 084006 (2010).
 - [39] <https://www.cascina.virgo.infn.it/advirgo/docs.html> (2012).
 - [40] B. Willke et al., Class. Quantum Grav. **23**, S207 (2006).
 - [41] T. Uchiyama et al., Class. Quantum Grav. **23**, S1161 (2004).
 - [42] Kuroda K for the LCGT Collaboration, Class. Quantum Grav. **27**, 084004 (2010).
 - [43] V. Mandic, E. Thrane, S. Giampanis, and T. Regimbau, Phys. Rev. Lett. (2012), forthcoming.
 - [44] N. Seto and K. Kyutoku, prd **86**, 042002 (2012).
 - [45] E. Thrane et al., Phys. Rev. D **80**, 122002 (2009).
 - [46] T. Regimbau, S. Giampanis, X. Siemens, and V. Mandic, Phys. Rev. D **85**, 066001 (2012).

- [47] P. A. Rosado, Phys. Rev. D **84**, 084004 (2011).
- [48] S. W. Ballmer, Class. Quantum Grav. **23**, S179 (2006).
- [49] S. Ballmer, Ph.D. thesis, Massachusetts Institute of Technology (2006).
- [50] S. Drasco and E. Flanagan, Phys. Rev. D **67**, 082003 (2003).
- [51] N. Seto, Phys. Rev. D **80**, 043003 (2009).
- [52] N. Seto, Astrophys. J. Lett. **683**, L95 (2012).
- [53] B. Allen and J. D. Romano, Phys. Rev. D **59**, 102001 (1999).
- [54] N. Christensen, Phys. Rev. D **46**, 4250 (1992).
- [55] E. Thrane, S. Kandhasamy, C. D. Ott, et al., Phys. Rev. D **83**, 083004 (2011).
- [56] T. Prestegard, E. Thrane, et al., Class. Quantum Grav. **28**, 095018 (2012).
- [57] Michael W Coughlin (for the LIGO Scientific Collaboration and the Virgo Collaboration), Class. Quantum Grav. **28**, 235008 (2011).
- [58] E. Thrane et al. (The Super-Kamiokande Collaboration), Astrophys. J. **704**, 503 (2009).
- [59] E. Thrane et al. (The Super-Kamiokande Collaboration), Astrophys. J. **697**, 730 (2009).
- [60] J. Braun et al., Astroparticle Physics **29**, 299 (2008).
- [61] B. Abbott et al., Phys. Rev. D **69**, 122004 (2004).
- [62] B. Abbott et al., Phys. Rev. Lett. **95**, 221101 (2005).
- [63] B. Abbott et al., Astrophys. J. **659**, 918 (2007).
- [64] C. Cutler and J. Harms, Phys. Rev. D **73**, 042001 (2006).
- [65] Typically, auto-power is calculated using adjacent time bins in order to avoid a bias that results from using the same data to calculate both $\hat{Y}(t; f)$ and $\hat{\sigma}(t; f)$
- [66] Using this pixel resolution, one year's worth of data would correspond to a significantly larger dataset— $\approx 3 \times 10^7$ measurements for every frequency bin—so this demonstration is not meant to be representative of a typical year-long stochastic search.